International Journal of Theoretical and Applied Finance
© World Scientific Publishing Company

arXiv:1010.3586v1 [stat.AP] 18 Oct 2010

A NONPARAMETRIC URN-BASED APPROACH TO INTERACTING FAILING SYSTEMS WITH AN APPLICATION TO CREDIT RISK MODELING

PASQUALE CIRILLO*
JÜRG HÜSLER
Institute of Mathematical Statistics and Actuarial Sciences, University of Bern
Sidlerstrasse 5, Bern, CH-3012, Switzerland

PIETRO MULIERE
Department of Decision Sciences, Bocconi University
Via Röntgen 1, Milan, IT-20136, Italy

Received (26^{th} May 2009)
Revised (17^{th} December 2009)

In this paper we propose a new nonparametric approach to interacting failing systems (FS), that is systems whose probability of failure is not negligible in a fixed time horizon, a typical example being firms and financial bonds.
The main purpose when studying a FS is to calculate the probability of default and the distribution of the number of failures that may occur during the observation period. A model used to study a failing system is defined *default model*.
In particular, we present a general recursive model constructed by the means of interacting urns.
After introducing the theoretical model and its properties we show a first application to credit risk modeling, showing how to assess the idiosyncratic probability of default of an obligor and the joint probability of failure of a set of obligors in a portfolio of risks, that are divided into reliability classes.

Keywords: Failing system; Urn model; Neutral to the right processes; Credit risk; Firms' defaults.

1. Introduction

A failing system (FS) is a system whose probability of failure is not negligible in a fixed time horizon (for example one year). The interest for such a topic is due to its diffusion in real-life problems: financial portfolios and credit risk, electrical and mechanical systems, firms' defaults, the world wide web can all be considered failing systems.
The main purpose when studying a FS is to calculate the probability of default and

*Corresponding author: pasquale.cirillo@stat.unibe.ch, tel. +41 (0)31 63 18 803.

the distribution of the number of failures that may occur during the observation period. A model that studies failing systems is defined *default model*.
Here we present new ways for calculating the probabilities of joint defaults in k different homogeneous groups of FS, when each group is characterized by some sort of external information about its reliability, which allows for an ordering. In particular, we propose a recursive model constructed by the means of interacting urns.
For simplicity, we assume that the probability of default is homogeneous within groups. Moreover we hypothesize that this probability of default is given by the sum of two components:

(1) a group idiosyncratic probability of failing that reveals, on the average, how a FS belonging to a given group is likely to fail "on its own";
(2) a systemic probability of failing, which represents the amount of negative/positive interactions among failing systems in different groups. In particular, once we have ordered the groups of FS from the best to the worst one, we will assume that the systemic probability of failing of group i increases (decreases) if the number of defaults in the superior groups $1, ..., i-1$ increases (decreases), while it remains the same for any change in the inferior groups $i+1, ..., k$.

Hence, the aim of this paper is to model the dependence among failures both within and between the k groups. This scheme can be efficiently reproduced with urns.
Urn processes (or urn models or urn schemes) constitute a very large family of probabilistic models in which the probability of certain events is represented in terms of sampling, replacing and adding balls in one or more urns or boxes.
Urn problems have been an important part of the theory of probability since the publication of the posthumous *Ars conjectandi* by Jakob Bernoulli [6] in 1713. Their most interesting characteristic is the possibility of simplifying complex probabilistic ideas, making them intuitive and concrete, and yet guaranteeing a good level of abstraction, that allows for general results.
The choice of urn processes as a probabilistic tool is mainly due to the following reasons:

(1) They are particularly suitable, thanks to their efficiency, for studying chance experiments, especially when these are characterized by countable spaces;
(2) They represent an excellent way to describe the concept of "random choice";
(3) Simple urns can be easily compounded into new ones in order to study more complex problems;
(4) Urn schemes have as powerful as elegant combinatorial properties, that allow for general, complex results in a rather concise form.
(5) There are many relationships and isomorphisms between urn models and other well-known mathematical objects (see for example *analytic urns* in [18]). All

this gives the possibility to the researcher of switching from one approach to the other at her/his convenience;
(6) Urns are very useful objects in simulations, given their natural connections with sampling schemes.

As discussed in [5] and [21], one of the prototypes of urn processes is the well-known Polya urn, developed at the beginning of the last century to model the diffusion of infectious diseases. It represents one of the simplest ways to generate beta exchangeable random variables and it is based on the concept of reinforcement. Moreover, its multidimensional version, as shown in [7], is a very useful tool to obtain the fundamental Dirichlet distribution, an essential tool for Bayesian statistics.
Polya urns represent one of the basic pillars of the present work. We will make use of Polya urns to define our urn chain model.
The analysis of default models via urns has several advantages:

(1) The modelization is rather intuitive and immediate;
(2) Urns can be considered a first attempt to study failing systems from a Bayesian nonparametric point of view, for they allow the researcher to introduce her/his prior knowledge into the analysis by modifying the initial urn composition and the reinforcement matrix of the urn process (see [8] for more details)
(3) The flexibility of urn schemes is a very useful characteristic for simulations and empirical studies, as we will see in the last section of the paper.

Some seminal ideas for our construction have been introduced in Marsili and Valleriani [22], where systems of interacting Polya urns are discussed. Anyway, the main references for our model, which is based on an iterative framework of interacting urns, are [15] and [26]. In particular, the basic ideas in [26], further developed in [9], can be considered the very starting point of this work.
The general framework we propose in this paper has, according to us, several interesting applications; a useful one being related to credit risk modeling.
Credit risk is the risk of loss due to a debtor's non-payment of a loan or other line of credit and it is strictly linked to the concept of default. One of the most important issues in credit risk modeling is represented by the assessment of the probability of failure/default of an obligor and/or a set of obligors in a portfolio of risks. In particular, especially for banks and financial companies, it is important to estimate the probability of joint defaults over a fixed time horizon, and this is why default models represent a fundamental tool in credit risk analysis. Furthermore, the assessment of the probability of joint defaults, as pointed out in [28], is also important for securities whose payoff is function of the profits and losses of a portfolio of underlying bonds. It follows that the study of dependence structure of interacting failing systems such as firms and bonds is fundamental for a correct estimation of credit risk, especially when the dependence cannot be summarized by simple measures of co-variability like linear correlation, despite this over-simplification is often used in credit risk modeling (see for example, [24]).

The paper is organized as follows: Section 2 introduces our model and the main probabilistic results; Section 3 propose a first study concerning credit risk modeling; and Section 4 concludes.

2. Introducing the urn chain model

We can now introduce the urn chain model in which several urns interact to reproduce dependent risks, the basic brick of our construction being represented by Polya urn.

Polya urn has been introduced by Polya and Eggenberger in 1923 [17] to study infectious and self-reinforcing phenomena. The most important characteristic of this particular urn scheme is its reinforcement mechanism, that has become the prototype for many probabilistic models for studying contagion and aftereffects. The behavior of Polya urn is very simple yet ingenious. In its simplest two-color version, imagine we have an urn containing balls of two different colors (say black and white). Every time we sample the urn we look at the color of the chosen ball and then put it back into the urn together with another ball of the same color. In this way, the more a given color has been sampled in the past, the more likely it will be sampled in the future. Obviously the reinforcement rule can be generalized introducing s balls of the same colors, considering a random or time-varying reinforcement and so on. For a complete analysis of Polya urn's combinatorics, behavior and generalizations see [21] and [23].

In what follows, we will make use of simple two-color Polya urns with general (but fixed) reinforcement s. Our choice of Polya urns to construct a risk model - with possible applications to credit risk - has been inspired by several works available in the literature, in which urns are used to model credit default distributions [3], allocation problems under uncertainty [4], actuarial problems [14], firms defaults [10], risk and ambiguity [19], just to cite some papers.

2.1. *The idiosyncratic probability of default*

Consider N failing systems (think of firms or bonds) divided into k groups that, for simplicity, we assume to be homogeneous. We also hypothesize that, within each group, the FS are exchangeable in the sense of de Finetti [1], that is to say that their joint probability is immune to permutations. This assumption is clearly weaker than the one of independence and identical distribution.

Each group consists of n_j elements, such that $\sum_{j=1}^{k} n_j = N$.

Assume that every group is characterized by some sort of external information (qualitative or quantitative) about its reliability, i.e. about the reliability of its components. In other words we ask every group to possess a score γ_j, $j = 1, ..., k$, such that the set $G = \{\gamma_j : j = 1, ..., k\}$ is a poset. This means that there exists a

relation $\precsim$ on G which is reflexive, antisymmetric and transitive or, formally,

$$\gamma_r \precsim \gamma_r$$
$$\text{if } \gamma_r \precsim \gamma_s \text{ and } \gamma_s \precsim \gamma_r \text{ then } \gamma_r = \gamma_s$$
$$\text{if } \gamma_r \precsim \gamma_s \text{ and } \gamma_s \precsim \gamma_t \text{ then } \gamma_r \precsim \gamma_t.$$

Without any loss of generality, we will assume that the k group are completely ordered according to their ratings and, specifically, $\gamma_1 \succsim \gamma_2 \succsim ... \succsim \gamma_{k-1} \succsim \gamma_k$. We will read the relation $\succsim$ as "better than" so, for example, $\gamma_1 \succsim \gamma_2$ means that group 1 is "better than" group 2 and, as a consequence of this, the $k-$th group is the worst one, since it is characterized by the lowest reliability. In particular, as signal of reliability we consider the idiosyncratic probability of default D_i of the different groups. In general - and this is neither a strong nor a ludicrous assumption - we want that $D_i < D_{i+1}$ for $i = 1, ..., k$.

As far as the idiosyncratic probability of default of every group is concerned, we want to construct a mechanism that updates that probability every time a failure occurs in a given group. As said, our idea is to use Polya urns (see [21]) that are characterized by a simple but efficient reinforcement rule.

So, let $D_i(t)$ represent the idiosyncratic probability of default of groups i at time t. Assume that every group is associated with a Polya urn U_i, $i = 1, ..., k$ initially containing $w_i(0) \geq 0$ white balls and $b_i(0) \geq 0$ black balls. At time $t = 1, 2, ...$, we sample a ball from urn U_i, look at its color and return it into the urn together with $s_i > 0$ additional balls of the same color. This mechanism, called reinforcement, evidently modifies the composition of the urn, updating the probability of picking a certain color. If we repeat the sampling infinite times, we obtain an infinite sequence of $0-1$ Bernoulli random variables $\{X_i(t)\}$, where $X_i(t) = 0$ if the sampled ball at time t is black and $X_i(t) = 1$ if white. The sequence $\{X_i(t)\}$ is exchangeable and it is called Polya sequence with parameters $(w_i(0), b_i(0), s_i)$.

For $t \geq 0$ let $W_i(t)$ and $B_i(t)$ represent the number of white and black balls in urn U_i at time t. It is easy to verify that

$$X_i(1) \sim Bern\left(\frac{w_i(0)}{w_i(0) + b_i(0)}\right) \tag{2.1}$$

and, in general,

$$X_i(t+1) \sim Bern\left(\frac{W_i(t)}{W_i(t) + B_i(t)}\right), \tag{2.2}$$

where $Bern(\eta)$ is the Bernoulli distribution of parameter η.

As far as the composition of the urn, the evolution rule is simply given by

$$(W_i(t+1), B_i(t+1)) = \begin{cases} (W_i(t) + s_i, B_i(t)) & \text{with probability } \frac{W_i(t)}{W_i(t)+B_i(t)}, \\ (W_i(t), B_i(t) + s_i) & \text{with probability } \frac{B_i(t)}{W_i(t)+B_i(t)}. \end{cases} \tag{2.3}$$

Proposition 2.1.
Let $\{X_i(t)\}$ be a Polya sequence with parameters $(w_i(0), b_i(0), s_i)$. Then:

(1) $\{X_i(t)\}$ *is exchangeable and its de Finetti measure is a Beta* $\left(\frac{w_i(0)}{s_i}, \frac{b_i(0)}{s_i}\right)$*;*

(2) the proportion of white balls $Z_i(t) = \frac{W_i(t)}{W_i(t)+B_i(t)}$ *converges with probability one to* p_i*.*

The proof of this proposition is very well-known and we refer to [21] for a complete demonstration. Here below we only sketch the basic arguments.

Proof.
First remember that $i = 1, ..., k$ are the different groups.
Point 1: let $1 \leq r \leq t$ and $(a_1, ..., a_t)$ such that $a_l \in \{0, 1\}$ and $\sum_{l=1}^{t} a_l = r$. Then

$$\begin{aligned} P(X_i(1) = a_1, ..., X_i(t) = a_t) &= \frac{\Gamma\left(\frac{w_i(0)}{s_1} + \frac{b_i(0)}{s_1}\right)}{\Gamma\left(\frac{w_i(0)}{s_i}\right)\Gamma\left(\frac{b_i(0)}{s_i}\right)} \frac{\Gamma\left(\frac{w_i(0)+r}{s_i}\right)\Gamma\left(\frac{b_i(0)+t-r}{s_i}\right)}{\Gamma\left(\frac{w_i(0)}{s_i} + \frac{b_i(0)}{s_i} + t\right)} \\ &= \int_0^1 \theta^r (1-\theta)^{t-r} \frac{\Gamma\left(\frac{w_i(0)}{s_1} + \frac{b_i(0)}{s_1}\right)}{\Gamma\left(\frac{w_i(0)}{s_i}\right)\Gamma\left(\frac{b_i(0)}{s_i}\right)} \theta^{\frac{w_i(0)}{s_i}-1}(1-\theta)^{\frac{b_i(0)}{s_i}-1} d\theta. \end{aligned} \tag{2.4}$$

For de Finetti's representation theorem this proves that the sequence is exchangeable. Furthermore, the unicity of the representation implies that the de Finetti measure of the sequence $\{X_i(t)\}$ is a *Beta* $\left(\frac{w_i(0)}{s_i}, \frac{b_i(0)}{s_i}\right)$.
Point 2: first we notice that $\{Z_i(t)\}$ is a bounded martingale, i.e.

$$\begin{aligned} E[Z_i(t+1)|Z_i(1), ..., Z_i(t)] &= \frac{W_i(t)+s_i}{W_i(t)+B_i(t)+s_i} \frac{W_i(t)}{W_i(t)+B_i(t)} \\ &\quad + \frac{W_i(t)}{W_i(t)+B_i(t)+s_i} \frac{B_i(t)}{W_i(t)+B_i(t)} \\ &= \frac{W_i(t)}{W_i(t)+B_i(t)} = Z_i(t) \end{aligned} \tag{2.5}$$

Hence, for Doob's convergence theorem, the sequence $\{Z_i(t)\}$ converges almost surely to a random limit $Z_i(\infty)$.
From Point 1 and the law of large numbers we know that, for t growing to infinity, the distribution of $t^{-1}\sum_{l=1}^{t} X_i(l)$ converges to a *Beta* $\left(\frac{w_i(0)}{s_i}, \frac{b_i(0)}{s_i}\right)$.
For every $t \geq 1$

$$Z_i(t) = \frac{w_i(0) + s_i \sum_{l=1}^{t} X_i(t)}{w_i(0) + b_i(0) + ts_i}, \tag{2.6}$$

and for every $z \in [0, 1]$

$$P(Z_i(t) \leq z) = P\left(t^{-1}\sum_{l=1}^{t} X_i(l) \leq z\left(\frac{w_i(0)}{ts_i} + \frac{b_i(0)}{ts_i} + 1\right) - \frac{w_i(0)}{ts_i}\right). \tag{2.7}$$

Hence

$$\lim_{t\to\infty} P(Z_i(t) \leq z) = \int_0^z \frac{\Gamma\left(\frac{w_i(0)}{s_i} + \frac{b_i(0)}{s_i}\right)}{\Gamma\left(\frac{w_i(0)}{s_i}\right)\Gamma\left(\frac{b_i(0)}{s_i}\right)} \theta^{\frac{w_i(0)}{s_i}-1}(1-\theta)^{\frac{b_i(0)}{s_i}-1} d\theta \tag{2.8}$$

□

Without any loss of generality (it is just a rescaling), from now on we assume $w_i(0) \in [0,1]$ and $b_1(0) = 1 - w_i(0)$.
Thus, before observing $\{X_i(t)\}$ the probability of picking a white ball is equal to $Z_i(0) = w_i(0)$. Anyway, as soon as n observation $X_i(1), ..., X_i(t)$ are available, we update our beliefs about the probability of sampling white balls, i.e.

$$Z_i(t) = \frac{w_i(0) + s_i \sum_{j=1}^{n} X_i(j)}{1 + n s_i}. \tag{2.9}$$

Equation 2.9 shows why the Polya urn is one of the basic tools in Bayesian nonparametrics. In fact, Equation 2.9 is consistent with the Bayesian paradigm of prior specification, knowledge update thanks to empirical observations and posterior calculation. Our prior knowledge is given by the urn composition at time 0. Then every time a white (or a black) ball is observed, our beliefs about the possibility of sampling white (or black) balls change and specifically increase, thanks to the reinforcement mechanism of the urn. Furthermore, given the initial composition and the updates, at every stage is possible to perform a prediction about the possibility of picking a given ball.
The informative contribution of every observation to the update process is given by the reinforcement quantity s_i, that is the number of balls added at every time step. It is clear that the relative contribution of an observation decreases with n, view that the more observation we have about process $\{X_i(t)\}$, the less we are ready to change our beliefs.
From Equation 2.9 and Proposition 2.1 we also know that $Z_i(t)$ and $R_i(t) = t^{-1}\sum_{l=1}^{t} X_i(l)$, the rate of ones over t, have the same limit. Hence, for t large enough, $R_i(t)$ is well approximated by the $Beta\left(\frac{w_i(0)}{s_i}, \frac{b_i(0)}{s_i}\right)$ distribution.
Coming back to the construction of our model, and remembering the Polya urn scheme, let us now assume that, for every group i, $E[D_i(0)] = E[Z_i(0)] = w_i(0)$. In other words, we want to associate the probability of default to the sampling of white balls in urn U_i; one default in group i corresponds to the extraction of one white ball from urn U_i. Moreover we make the hypothesis that $D_i(0)$ is distributed according to a $Beta(\frac{w_i(0)}{s_i}, \frac{1-w_i(0)}{s_i})$.
Given $D_i(t)$, the default of element j in group i over the time horizon t is distributed as a $Bern(D_i(t))$. If n_i is the number of elements in group i, and for $j = 1, ..., n_i$, we let $\delta_i^j(t)$ represent the indicator function of the event "default of element j in group i up to time t". In other words $\delta_i^j(t) = 1$ if element j has defaulted at some point in the time interval [0,t] and $\delta_i^j(t) = 0$ otherwise. We also let $\delta_i^1(t), \delta_i^2(t), ..., \delta_i^{n_i}(t)$ be conditionally independent given $D_i(t)$.

Without any empirical observation and following our a priori knowledge, the probability of default at time t, $D_i(t)$,is clearly equal to $w_i(0)$. Anyway, it is very likely that, after a time t has passed, we perfectly know how many elements of group i have failed, thus we are ready to update our beliefs. In particular, following the Polya urn mechanism, we add s_i white balls for every observed default and s_i black balls for every surviving element. As a consequence of this $D_i(t)$ is distributed as a *Beta* whose parameters are:

$$w_i^*(t) = \frac{w_i(0) + s_i \sum_{j=1}^{n_i} \delta_i^j(t)}{s_i} \quad \text{and} \quad b_i^*(t) = \frac{(1 - w_i(0)) + s_i \left(n_i - \sum_{j=1}^{n_i} \delta_i^j(t)\right)}{s_i}. \tag{2.10}$$

Moreover, thanks to Bayes theorem, it is straightforward to verify that the conditional distribution of $D_i(t)$ given $\delta_i^1(t), \delta_i^2(t), ..., \delta_i^{n_i}(t)$ is a *Beta* whose parameters are also expressed in Equation 2.10, and that

$$E[D_i(t)|\delta_i^1(t), \delta_i^2(t), ..., \delta_i^{n_i}(t)] = \frac{w_i(0) + s_i \sum_{j=1}^{n_i} \delta_i^j(t)}{1 + s_i n_i}. \tag{2.11}$$

Thanks to this simple Polya-like urn scheme, we have thus modeled the idiosyncratic probability of default for every group. It is easy to see a clear relationship between the standard Polya urn and our adaptation to the idiosyncratic probability of default in group i, once we notice that, by forcing a little bit the notation, $X_i(t) = \sum_{j=1}^{n_i} \delta_i^j$. Finally, it is important to notice that, being $w_i(0)$, $b_i(0)$ and s_i generally different among groups, we are dealing with distinct *Beta* distributions. Our use of the *Beta* distributions for the idiosyncratic probabilities of default is consistent with several empirical and theoretical studies, as underlined in [3].

2.2. *Modeling interaction: the systemic probability of default*

In our construction, the systemic probability of default accounts for the dependence among groups. In particular, once we know the idiosyncratic probabilities of default of the k groups and we have ordered them from the most reliable to the least reliable one, we want the superior/best groups to a have a direct influence on the inferior/worst ones.

Let once again $D_1(t), D_2(t), ..., D_k(t)$ be the idiosyncratic probabilities of default for the k groups. It is evident that $D_i(t) \in (0, 1)$. Now, define $D_i^*(t)$ as the total probability of default associated to group i at time t, that's the "sum" of the idiosyncratic and the systemic components. Having in mind the construction of [15] and [26] for neutral to the right processes, [20] for stick-breaking priors, and [9] for default models, we construct the probabilities of failure of the k groups as

$$\begin{aligned} D_1^*(t) &= D_1(t) \\ D_2^*(t) &= D_1^*(t) + (1 - D_1^*(t))D_2(t) \\ &\vdots \\ D_k^*(t) = D_{k-1}^*(t) + (1 - D_{k-1}^*(t))D_k(t) &= 1 - \prod_{i=1}^{k} (1 - D_i(t)) . \end{aligned} \tag{2.12}$$

It is easy to verify that:

(1) this construction respects all our assumptions, so that the better groups of FS's show a lower probability of default;
(2) the probabilities of default of the different groups are strictly linked together by the means of the recursive scheme.

So, thanks to this simple and rather intuitive iterative modeling, we obtain the probabilities of default for the different risk groups and, for every FS, we are able to say whether it is likely to fail.
Please note that from now on we are omitting time t not to perplex the notation; in other words, $D_i^* = D_i^*(t)$ and so on.

Proposition 2.2.
The process that governs the probabilities of failure D_i^*, $i = 1, ..., k$, is a neutral to the right process.

Proof.
First of all let $E_i = D_i^* - D_{i-1}^*$ with $i = 1, ..., k$ and $E_1 = D_1^* = D_1$. It is easy to verify that

$$(E_1, E_2, ..., E_k) \stackrel{d}{=} \left(D_1, D_2 (1 - D_1), ..., D_k \prod_{i=1}^{k-1} (1 - D_i) \right). \tag{2.13}$$

This, as shown in remark 3.1*b* in Doksum [15], assures that the process governing $(D_1^*, D_2^*, ..., D_k^*)$ is neutral to the right. □

Neutral to the right processes have been introduced by [15] and are widely used in Bayesian nonparametrics for survival analysis. For a complete introduction to this type of processes, we refer to the original paper by Doksum [15] and to other more recent works like [26] and [29]. Here it is sufficient to state the following definition.

Definition 2.1 (Doksum [15]).
The random distribution function F is said to be neutral to the right if for each $h > 1$ and $t_1 < ... < t_h$, there exist nonnegative independent random variables $V_1, ..., V_h$ such that

$$(F(t_1), F(t_2), ..., F(t_h)) =_{\mathcal{L}} \left(V_1, 1 - (1 - V_1)(1 - V_2), ..., 1 - \prod_{i=1}^{h} (1 - V_i) \right).$$

The equations

$$F(t_j) = 1 - \prod_{i=1}^{j} (1 - V_i) \quad j = 1, ..., h$$

yield

$$F(t_j) - F(t_{j-1}) = V_j \prod_{i=1}^{j-1} (1 - V_i)$$

and

$$V_j = \frac{(F(t_j) - F(t_{j-1}))}{(1 - F(t_{j-1}))} \quad j = 1, ..., h \text{ and } t_0 = -\infty.$$

Thus "F is neutral to the right" mainly means that the normalized increments

$$F(t_1), \frac{(F(t_2) - F(t_1))}{(1 - F(t_1))}, ..., \frac{(F(t_h) - F(t_{h-1}))}{(1 - F(t_{h-1}))}$$

are independent for all the $t_1 < ... < t_h$.

The fact that we have modeled the idiosyncratic probabilities of default by the means of Polya urns has an interesting consequence.

Proposition 2.3.
In particular the process that governs the probabilities of failure D_i^*, $i = 1, ..., k$, is a beta-Stacy process with parameters $(w_1^*, b_1^*; w_2^*, b_2^*; ...; w_k^*, b_k^*)$, where $w_i^* = w_i^*(t)$ and $b_i^* = b_i^*(t)$ are defined as in equation 2.10.

Proof.
Since the beta-Stacy process, as defined in [29], is a special case of neutral to the right process, when the independent variables are Beta distributed, it is straightforward to prove the proposition. In fact, using equation 2.13 we known that

$$\begin{aligned} E_1 &= D_1 \\ E_2 &= D_2(1 - D_1) \\ &\vdots \\ E_k &= D_k \prod_{i=1}^{k-1} (1 - D_i). \end{aligned}$$

An obvious consequence of this is that

$$\begin{aligned} E_1 &\sim \mathcal{BS}(w_1^*; b_1^*; 1) \\ E_2|E_1 &\sim \mathcal{BS}(w_2^*; b_2^*; 1 - E_1) \\ &\vdots \\ E_k|E_{k-1}, ..., E_1 &\sim \mathcal{BS}(w_k^*; b_k^*; 1 - \sum_{j=1}^{k-1} E_j), \end{aligned}$$

where $\mathcal{BS}(a;b;c)$ is the so-called beta-Stacy distribution introduced by [25], whose density function is

$$\frac{1}{B(a,b)}x^{a-1}\frac{(c-x)^{b-1}}{c^{a+b-1}}I_{(0,c)}(x),$$

with $B(a,b)$ representing the standard beta function.
Hence the final result immediately follows. □

We would like to stress that, while proposition 2.3 is strictly linked to the use of Polya urns and *Beta* distributions to model the idiosyncratic probabilities of defaults, proposition 2.2 holds in general. The only simple requirement is that the variables D_i are i.i.d. and $D_i \in [0,1]$ for $i = 1, ..., k$. In other words, given the recursive construction, the urn chain, independently from the type of urn used, always generates neutral to the right processes. This result is quite useful in practice, since neutral to the right processes are conjugate, thus simplifying Bayesian prediction.

Corollary 2.1.
If $b_i = \sum_{k>i} w_i$ *(and* $\sum_{i=1}^{\infty} w_i < \infty$*) the beta-Stacy process that governs the probabilities of default simply becomes a Generalized Dirichlet distribution,* $GD\left(w_1^*, b_1^*; w_2^*, b_2^*; ...; w_k^*, b_k^*\right)$*, as defined in [11].*

Proof.
The proof is an application of Theorem 4.1 in Muliere and Walker [26]. □

This last corollary, whose conditions are easily fulfilled, is very useful in applications. In fact, if the beta-Stacy process degenerates to a Generalized Dirichlet distribution, its parameters can be easily estimated using several existing computational techniques, from the expectation-maximization algorithm to the generalized method of moments and other more advanced tools (see for example []).
Now, let F_i, $i = 1, ..., k$, be the number of failures in the $i-$th group with n_i elements. The (marginal) probability of having $F_i = f_i$ failures in the $i-$th group, with $0 \leq f_i \leq n_i$, is equal to

$$P\left[F_i = f_i\right] = E\left[\binom{n_i}{f_i}(D_i^*)^{f_i}(1-D_i^*)^{n_i-f_i}\right] = E\left[P\left[F_i = f_i | D_{i-1}^*\right]\right]. \quad (2.14)$$

Then, using standard combinatorial considerations, the joint defaults of the first two groups can be computed as

$$P\left[F_1 = f_1, ..., F_k = f_k\right] = E\left[\binom{n_1}{f_1}(D_1^*)^{f_1}(1-D_1^*)^{n_1-f_1}\binom{n_2}{f_2}(D_2^*)^{f_2}(1-D_2^*)^{n_2-f_2}\cdots \right.$$
$$\left. \cdots \binom{n_k}{f_k}(D_k^*)^{f_k}(1-D_k^*)^{n_k-f_k}\right] \quad (2.15)$$

Given these details we have that, at time t, the number of default in the first group (the best one) follows a beta-binomial distribution (see [7]), or

$$P[F_1(t) = f_1] = \binom{n_1(t)}{f_1}\frac{B(w_1^*(t), b_1^*(t))}{B(w_1(0), b_1(0))}, \quad (2.16)$$

where $n_1(t)$ is the number of failing systems in group 1 at time t. This results comes directly from the use of Polya urns.
At this point, we can obtain the number of defaults in the first two groups, that is

$$
\begin{aligned}
&P[F_1(t) = f_1, F_2(t) = f_2] = E[P[F_1(t) = f_1, F_2(t) = f_2 | D_1^*(t), D_2^*(t)]] \\
&= \binom{n_1(t)}{f_1}\binom{n_2(t)}{f_2} E[D_1^*(t)^{f_1} D_2^*(t)^{f_2} (1 - D_1^*(t))^{n_1(t)-f_1} (1 - D_2^*(t))^{n_2(t)-f_2}] \\
&= \binom{n_1(t)}{f_1}\binom{n_2(t)}{f_2} E\left[D_1(t)^{f_1}(1 - D_1(t))^{n_1(t)-f_1}(1 - D_1(t))^{n_2(t)-f_2}(1 - D_2(t))^{n_2(t)-f_2} \times\right. \\
&\left.\times \left(\sum_{i=0}^{f_2}\binom{f_2}{i} D_1(t)^i D_2(t)^{n_2(t)-i}(1 - D_1(t))^{n_2(t)-i}\right)\right] = \binom{n_1(t)}{f_1}\binom{n_2(t)}{f_2}\times \\
&\times \sum_{i=0}^{f_2}\binom{f_2}{i}\frac{B(w_1^*(t) + i, b_1^*(t) + n_2(t) - i)}{B(w_1(0), b_1(0))}\frac{B(w_2^*(t) - i, b_2^*(t))}{B(w_2(0), b_2(0))}.
\end{aligned}
\tag{2.17}
$$

It should be now clear that the joint probability of the number of defaults in the k groups can be obtained continuing the iterative construction of equations 2.16 and 2.17 and standard combinatorial techniques, the result being a combination of beta-binomial distributions.

3. An application to credit risk modeling

The way in which international rating companies such as Moody's, S&P and Fitch deal with credit risk and firms' defaults is definitely similar to the framework of interacting failing systems we have introduced. Think for example of N firms divided into k homogenous groups, that are ordered according to their financial reliability. For these reasons, we here present a simulation exercise related to firms' defaults and credit risk.
In order to use our model in applications, we need at least to know the quantities $D_i(0)$ for $i = 1, ..., k$, that is the idiosyncratic probabilities of default for every group at time 0. In general this should not be a problem, since the model can be initialized with historical data.
Here we propose a first application of our model to firms' defaults. In particular we show a simulation exercise using fictitious data. The suitability of the urn chain model for this kind of phenomena is definitely supported by empirical evidence (see [2]).
Imagine we have 290 firms divided into three groups of reliability: A, B, C. Group A contains the 20 best firms on the market, whose probability of default is very low. Group C is made up of the riskiest firms and has 180 elements. Group B is the intermediate one and contains 90 firms. Even though simplified, this framework correctly reproduces the firms' classification structure induced by rating companies such as Moody's and Standard & Poor's.

To initialize the model, i.e. to define the values $D_i(0)$ for $i = A, B, C$, we can use different methods. A first possibility is given by historical data: looking at the time series of defaults, we can try to estimate the three idiosyncratic probabilities of default and then run the urn chain. Another possibility is fully Bayesian and it is related to the prior knowledge of the researcher. In fact it is always possible to incorporate one's beliefs in $D_i(0)$ and then use the urn scheme as updating mechanism (see [8]). Finally, a third way for setting the model up is given by the credit spread approach of [16], which is the solution we adopt here.

Following [16], we consider that the value of the whole probability of default $D_i^*(1)$ for period $[0, 1]$ is indirectly quoted on the market as the average one-year credit spread $\gamma_i(1)$ of group i. In other words, $\gamma_i(1)$ measures the average riskiness of a member of group i as the difference between the zero coupon bond of that member and the risk-free interest rate on the market. The values of $\gamma_i(1)$ can be quite easily found on the market and are surely available to practitioners. In particular, in the financial literature (see [3]) it is common to assume

$$E[D_i^*(t)] = 1 - \exp(-t\gamma_i(t)). \tag{3.1}$$

Hence $E[D_i^*(1)] = 1 - \exp(-\gamma_i(1))$.

Since the quantity we are interested is $E[D_i(0)] = w_i(0)$, that is the idiosyncratic probability of default at time 0, we can try to obtain it from $D_i^*(1)$, using a linear term structure for credit spreads commonly used by traders (see for example [13]). In fact, while it is not a problem to obtain actual 1-year default probabilities, it can be harder to have estimates for shorter periods. In general, traders do assume that, once we have fixed a time horizon, the risk of having one or more defaults before the expiry linearly decreases.

The use of an underlying linear term structure can be also a good way for making predictions about the 1-year probability of default every time some new information is available about the numbers of defaults in the different groups. In what follows we split every year in twelve months, but the same reasoning is available for weeks and even days.

In our construction we have decided to order the groups of failing systems using their idiosyncratic probability of default as rating measure. Obviously there are many possibilities for ordering and they all depend on the amount of available information. In our case, assuming that $\gamma_A(1) = 0.02$, $\gamma_B(1) = 0.06$ and $\gamma_C(1) = 0.09$ are the one-year credit spreads for the three groups, and looking at the idiosyncratic probability of default we have that

$$\begin{aligned} &E[D_A^*(1)] = E[D_A(1)] = 1 - \exp(-\gamma_A(1)) = 0.0198 \\ &E[D_B^*(1)] = E[D_A^*(1) + (1 - D_A^*(1))D_B(1)] = 1 - \exp(-\gamma_B(1)) = 0.0582 \\ &E[D_C^*(1)] = E[D_B^*(1) + (1 - D_B^*(1))D_C(1)] = 1 - \exp(-\gamma_C(1)) = 0.0861. \end{aligned} \tag{3.2}$$

Thanks to equation 3.1 we know that, for group A - the best one, the total probability of default is equal to the idiosyncratic component. For the other two groups, on the contrary, the total probability of default also includes a systemic component

that accounts for the dependence between groups. As a consequence of this we have that

$$E[D_B(1)|D_A^*(1)] = \frac{1 - \exp(-\gamma_B(1)) - D_A^*(1)}{1 - D_A^*(1)}, \tag{3.3}$$

and

$$E[D_C(1)|D_B^*(1)] = \frac{1 - \exp(-\gamma_C(1)) - D_B^*(1)}{1 - D_B^*(1)}. \tag{3.4}$$

Now let us hypothesize that, according to the linear term structure, the credit spread decreases of 0.0005 points every month. This value is fictitious but it is not very far from what practitioners generally assume (see [16] and [27]). In this way, if $\gamma_A(1) = 0.02$ then $\gamma_A(6/12) = 0.02 + 0.0005 * 6 = 0.0230$, $\gamma_A(3/12) = 0.0245$ and so on. Hence we have $\gamma_A(0) = 0.0260$, $\gamma_B(0) = 0.0660$ and $\gamma_C(0) = 0.0960$.
Trivially for periods less than one year we have that for group A (and the same holds for B and C) the probability of failing between time 0 and $i/12$ is such that

$$E[D_A^*(i/12)] = 1 - \exp(-\frac{i}{12}\gamma_A(i/12)). \tag{3.5}$$

Now let us assume that for every group, at the end of every month, we observe the number of defaults expressed in the tables 1, 2 and 3. For example in the first month we have 0 defaults for group A, 3 in group B and 25 in C.
We finally need to define the values of reinforcement s_i for $i = A, B, C$. For simplicity we assume that the reinforcement is always the same for all the groups and equal to 0.05 and 0.01. Obviously one can define diverse reinforcement rules for the different groups.
Since s_i represents the size of information update generated by every observation, we understand that a greater value is equivalent to a considerable reinforcement, while 0.01 corresponds to a weaker update. Probably, as we show at the end of this section, the choice of s_i is one of the most sensible features of our model: from one side, it allows the researcher to incorporate an eventual a priori knowledge about the impacts of defaults; from the other, different values can produce quite different results in estimation (see tables 1, 2 and 3). A good idea could be to calibrate s_i such that the variability of the reinforced credit spreads is as close as possible to the historical variability of the spreads quoted on the market.
Using equations 2.10, 2.11 and 3.1 is now possible to perform our simulation, obtaining the results of tables 1, 2 and 3 (notice that "uc" stands for urn chain).

Every table is devoted to one of the three groups A, B, C. Each table contains the following information: the months from 0 (present time) to 12 (one year); the values of the credit spread at one year λ_i in the different months according to the linear spread term structure; the number of defaults in the different months; the estimated probability of default according to our urn chain model with two different values for the reinforcement quantity s_i, that is 0.01 and 0.05.
As expected our model gives estimates of the probability of default that are clearly different form the basic predictor based on equation 3.1. In particular, it is evident

Table 1. Simulation results for group A

Month (i)	0	1	2	3	4	5	6
$\gamma(i/12)$	0.0260	0.0255	0.0250	0.0245	0.0240	0.0235	0.0230
Defaults	0	0	1	0	0	0	2
$E[D^*_{uc,0.05}(i/12)]$	0.0257	0.0128	0.0314	0.0161	0.0083	0.0042	0.0535
$E[D^*_{uc,0.01}(i/12)]$	0.0257	0.0214	0.0262	0.0220	0.0185	0.0155	0.0298
Month (i)	7	8	9	10	11	12	
$\gamma(i/12)$	0.0225	0.0220	0.0215	0.0210	0.0205	0.0200	
Defaults	1	0	0	0	0	0	
$E[D^*_{uc,0.05}(i/12)]$	0.0559	0.0311	0.0173	0.0096	0.0053	0.0030	
$E[D^*_{uc,0.01}(i/12)]$	0.0341	0.0294	0.0253	0.0218	0.0188	0.0162	

Table 2. Simulation results for group B

Month (i)	0	1	2	3	4	5	6
$\gamma(i/12)$	0.0660	0.0655	0.0650	0.0645	0.0640	0.0635	0.0630
Defaults	0	3	1	0	4	5	8
$E[D^*_{uc,0.05}(i/12)]$	0.0639	0.0468	0.0467	0.0190	0.0462	0.0605	0.1426
$E[D^*_{uc,0.01}(i/12)]$	0.0639	0.0570	0.0503	0.0350	0.0466	0.0581	0.0974
Month (i)	7	8	9	10	11	12	
$\gamma(i/12)$	0.0625	0.0620	0.0615	0.0610	0.0605	0.0600	
Defaults	9	5	5	4	0	2	
$E[D^*_{uc,0.05}(i/12)]$	0.1714	0.1212	0.1072	0.0921	0.0304	0.0408	
$E[D^*_{uc,0.01}(i/12)]$	0.1253	0.1170	0.1135	0.1069	0.0773	0.0698	

how the numbers of defaults in the different months have a clear impact on the probability of default.

This information about defaults would be probably neglected in a standard approach without reinforcement or, in the best case, all the update would be performed at the end of the 1-year period, when all defaults have happened, and only as a basis for the next period of interest, with a clear temporal delay. Unfortunately this eventual all-in-one update is not really useful at all, since it can make the probability of default increase for all the second year, even if during the second year no default actually happens. In other words, a practitioner could base his/her evaluations on the basis of out-of-date facts. Our model is instead continuously updated, always producing updated estimates.

For every group, the role of the urn reinforcement mechanism is really clear if we

Table 3. Simulation results for group C

Month (i)	0	1	2	3	4	5	6
$\gamma(i/12)$	0.0960	0.0955	0.0950	0.0945	0.0940	0.0935	0.0930
Defaults	0	25	19	9	14	10	24
$E[D^*_{uc,0.05}(i/12)]$	0.0915	0.1688	0.1641	0.0911	0.1466	0.1460	0.3225
$E[D^*_{uc,0.01}(i/12)]$	0.0915	0.1512	0.1583	0.1183	0.1417	0.1464	0.2458
Month (i)	7	8	9	10	11	12	
$\gamma(i/12)$	0.0925	0.0920	0.0915	0.0910	0.0905	0.090	
Defaults	15	14	9	9	9	7	
$E[D^*_{uc,0.05}(i/12)]$	0.3321	0.3083	0.2763	0.2824	0.2764	0.3101	
$E[D^*_{uc,0.01}(i/12)]$	0.2789	0.2869	0.2805	0.2832	0.2782	0.2874	

compare the expected probability of default with $s_i = 0.05$ and $s_i = 0.01$. The greater is the reinforcement the greater are the fluctuations of the expected probability of default after every update. For example $s_i = 0.05$ seems to be a quite high value for the updating process: if we consider group A and periods 5 and 6, we see that with $s_A = 0.05$ the probability of default jumps from 0.0042 to 0.0535, indicating an excessive sensitivity of the model to defaults. For $s_A = 0.01$, on the contrary, the jump is more contained, from 0.0155 to 0.0298, suggesting a more plausible variation.

Comparing the three tables, we can finally notice that, thanks to the urn chain mechanism, every time the probability of default increases (decrease) in group A, the dependence structure makes the probabilities of default of the inferior groups increase too. In other words, avoiding the assumption of independent defaults (both within and between groups), we have tried to overcome one of the weakest points of standard credit risk models (see for example the CR+ model in [12]).

4. Conclusion

In this paper we have shown a first (Bayesian) nonparametric model for studying failing systems.

In detail, we consider failing systems divided into homogeneous groups of different reliability and we assume that these groups can be ordered according to some sort of external information. The elements within each group are assumed to be exchangeable.

We hypothesize that the probability of default of every failing system is given by the sum of two different components: an idiosyncratic probability of default related to the group to which the FS belongs and a systemic probability of defaults that account for the dependence between groups. For the first probability we make use of

Polya urns that allows for a Bayesian nonparametric modeling based on information update, while for the second probability we construct an urn chain whose structure is very close to the one used by [15] for constructing neutral to the right processes. We have also proposed a possible application of our model; in particular, we show a simulation experiment related to credit risk modeling. At this point it could be worth to apply the model to actual problems and data and to compare it with some benchmark.

As far as the evolution of the model is concerned it could be interesting to substitute the simple Polya mechanism with a more advanced scheme. An idea could be to use reinforced urn processes (see [29]) to model the idiosyncratic probabilities of defaults and then to combine them using the same neutral to the right construction. We believe that the general process governing the probability of default would still be a beta-Stacy process, but further analysis is needed.

Another research line could be to introduce the possibility of transitions from one group to the other, in order to model down- and upgrading of firms in reliability classes. This could be done by introducing random reinforcement rules.

Finally, we would like to stress that the choice of Polya urns to model the idiosyncratic probability of default is only due to the desire of obtaining closed-form results for the number of defaults. In reality, several other urn schemes could be used, while maintaining the general neutral to the right structure of the model (see [9] for more details).

Acknowledgements: The authors would like to thank the editor and the referees for their valuable suggestions. This work has been partly supported by the Swiss National Science Foundation, to which the first two authors are particularly grateful.

References

[1] J. Aldous. *Exchangeability and related topics*. Lecture Notes in Mathematics, Springer (1985).

[2] L. Allen, Credit risk modelling of middle markets, *School of Business Scientific Report* **11** (2004).

[3] E. Amerio, P. Muliere and P. Secchi, Reinforced urn processes for modelling credit default distribution, *International Journal of Theoretical and Applied Finance* **7** (2004), 407-423.

[4] W.B. Arthur, Competing Technologies, Increasing Returns, and Lock-In by Historical Events, *The Economic Journal*, **394** (1989), 116-131.

[5] N. Balakrishnan (Ed.), *Advances in combinatorial methods and applications to probability and statistics*, Birkhauser (1997).

[6] J. Bernoulli, *Ars Conjectandi*, 1713. English version in Dudley Sylla E., editor, *The Art of Conjecturing*, JHU Press (2006).

[7] D. Blackwell and J.B. MacQueen, Ferguson distribution via Polya urn schemes, *Annals of Statistics* **1** (1977), 353-355.

[8] D.M. Cifarelli, P. Muliere and P. Secchi, Urn schemes for constructing priors. *Technical Report 429/P*, Dipartimento di Matematica Politecnico di Milano, (2000).

[9] P. Cirillo, New urn approaches to shock and default models, *PhD Thesis*, Bocconi University, (2008).
[10] P. Cirillo and J. Hüsler, Shock models for defaults: parametric and nonparametric approaches, to appear in the *Festschrift in honor of Tom Hettmansperger*, WSP, (2010).
[11] R.J. Connor and J.E. Mosiman, Concepts of independence for proportions with a generalization of the Dirichlet distribution, *Journal of the American Statistical Association* **64** (1969), 194-206.
[12] Credit Suisse Financial Products. CreditRisk+: a credit risk management framework, *Technical Report* available online, (1997).
[13] Q. Dai and K. Singleton, Term structure dynamics in theory and practice, *Review of Financial Studies* **16** (2003), 631-678.
[14] C.D. Daykin, T. Pentikinen and M. Pesonen, *Practical risk theory for actuaries*, CRC Press, (1994).
[15] K. Doksum, Tailfree and neutral random probabilities and their posterior distributions, *Annals of Statistics* **2** (1974), 183-201.
[16] D. Duffie and K.J. Singleton, *Credit Risk*, Princeton University Press, (2003).
[17] F. Eggenberger and G. Polya. Über die statistik verketteter Vorgäge. *Zeitschrift für Angewandte Mathematik und Mechanik* **3**, 279-289.
[18] P. Flajolet, J. Gabarro and H. Pekari, Analytic urns, *Annals of Probability* **33** (2005), 1200-1233.
[19] D. Ghosh and M. Ray, Risk, Ambiguity, and Decision Choice: Some Additional Evidence,*Decision Sciences* **28** (1997), 81-104.
[20] H. Ishwaran and L.F. James, Gibbs sampling methods for stick-breaking priors, *Journal of the American statistical Association* **96** (2001), 161-173.
[21] N.L. Johnson and S. Kotz, *Urn models and their applications*, Wiley-Interscience, (1977).
[22] M. Marsili and A. Valleriani, Self-organization of interacting Polya urns, *Preprint* at citeseer.ist.psu.edu/194060.html, (1998).
[23] H.M. Mahmoud, *Polya Urn Models*, CRC Press, (2009).
[24] A.J. McNeil, R. Frey and P. Embrechts, *Quantitative Risk Management*, Princeton Press, (2005).
[25] G.A. Mihram and R.A. Hultquist, A bivariate warning-time/failure distribution, *Journal of the American Statistical Association* **62** (1967), 589-599.
[26] P. Muliere and S. Walker, Neutral to the right processes from a predictive perspective: a review and new developments, *Metron* **58** (2000), 13-30.
[27] P. Schönbucher, The term structure of defaultable bond prices, *Review of Derivatives Research* **2** (1998) 161-192.
[28] P. Schönbucher, *Credit Derivatives Pricing Models: Models, Pricing and Implementation*, Wiley, (2003).
[29] S. Walker and P. Muliere, Beta-stacy processes and a generalization of the Polya urn scheme, *Annals of Statistics* **25** (1997), 1762-1780.